Lectures in Mathematical Modeling

Lecture 3: Model Fitting and Residual Analysis

Kenneth Sok Kin Cheng

Revised 20 August 2026

Learning goals

By the end of the lesson, readers should be able to:

1. distinguish **model fitting** from **interpolation**;
2. calculate and interpret a **residual**;
3. fit a proportional model and a straight-line model by the **least-squares method**;
4. linearise simple exponential or power relationships using transformed variables;
5. compare fitted models using graphs, residual patterns, error measures, and withheld data;
6. communicate why a numerically good fit may still be an unsuitable model.

1 From a proposed relationship to a fitted model

In Lecture 2, physical reasoning suggested relationships such as

$$d_b \propto v^2$$

for braking distance and speed. Such a statement describes the *form* of a model, but it does not yet determine the proportionality constant. Real observations are also unlikely to lie exactly on a mathematical curve. We therefore need a systematic way to estimate parameters and judge whether the proposed relationship is credible.

Launch: Three lines, three stories **[10 min]**

A set of six data points is shown with three possible fitted lines drawn by different teams.
In pairs, discuss:

1. What could "best fit" mean?
2. Should a positive error cancel a negative error?
3. Should one very inaccurate point matter more than several small inaccuracies?
4. What information would you need before choosing among the lines?

Modeling lesson: A fitted model is never selected by appearance or a single number alone. The criterion must match the purpose of the model.

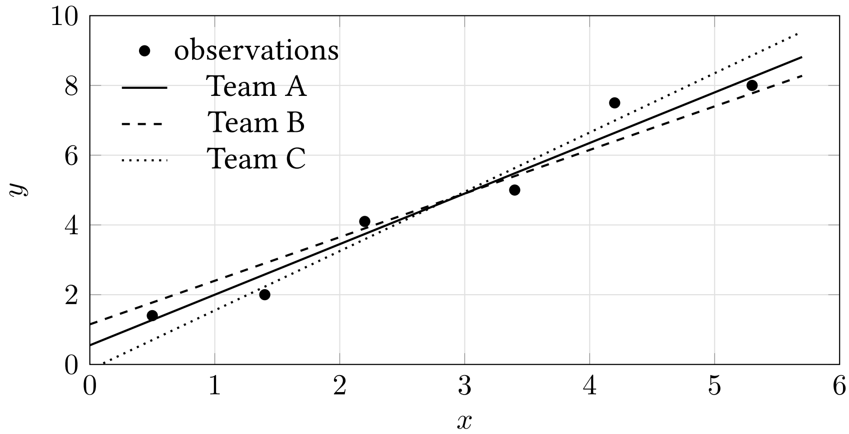

Figure 1: Different fitting criteria can prefer different models. The purpose of fitting must be stated.

Fitting, interpolation, and prediction

- In **model fitting**, a model family is proposed from theory or previous reasoning, and its parameters are estimated from data. The curve need not pass through every point.
- In **interpolation**, a curve is constructed to pass through specified data points, usually to estimate values between them. This is the main subject of Lecture 4.
- A fitted model may be used for **prediction**, but prediction outside the observed range is **extrapolation** and requires additional caution.

2 Residuals and sources of discrepancy

Suppose the observed value at x_i is y_i and the model predicts $\hat{y}_i = f(x_i)$. The **residual** is

$$\boxed{r_i = y_i - \hat{y}_i.}$$

A positive residual means the model underpredicts the observation; a negative residual means the model overpredicts it.

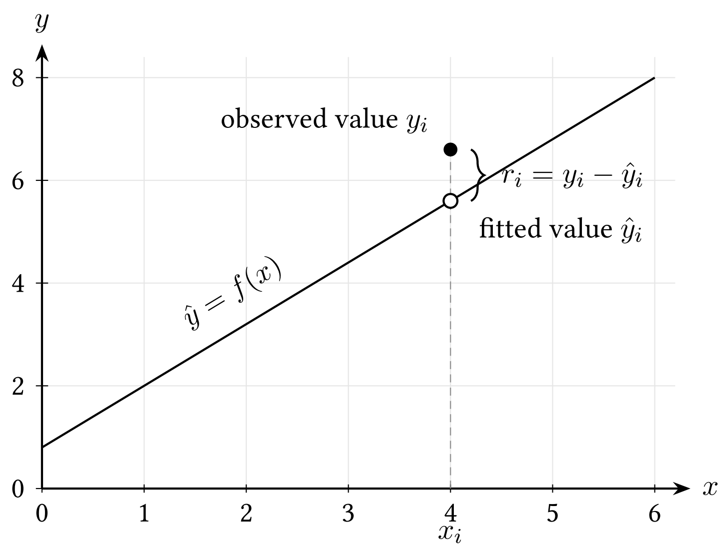

Figure 2: At x_i, the residual is the signed vertical difference between the observed value y_i and the fitted value $\hat{y}_i = f(x_i)$.

A discrepancy is not always a fitting failure

Differences between data and model can arise from several sources:

1. **measurement error**: limited precision, recording error, or natural variation;
2. **formulation error**: omitted variables or an unsuitable functional form;
3. **parameter uncertainty**: the available data do not determine the parameters exactly;
4. **numerical error**: rounding or approximation during computation.

A very precise fitting algorithm cannot repair a poor model structure or unreliable data.

Checkpoint: Signs and units

A model predicts a stopping distance of 42.3 m, while the observed distance is 45.0 m.

1. Calculate the residual.
2. Does the model overpredict or underpredict?
3. What are the units of the residual?

3 Fitting a proportional model through the origin

If theory suggests

$$y = kx,$$

we should not automatically fit a general line $y = ax + b$. A nonzero intercept would contradict the proposed mechanism. For data (x_i, y_i), the least-squares value of k minimises

$$S(k) = \sum_{i=1}^{m} (y_i - kx_i)^2.$$

Differentiating and setting $S'(k) = 0$ gives

$$\boxed{k = \frac{\sum x_i y_i}{\sum x_i^2}.}$$

Worked Example 1: Calibrating the braking-distance model

Lecture 2 proposed $d_b = kv^2$, where v is speed in m/s and d_b is braking distance in m. Let

$$X = v^2.$$

Then the model becomes $d_b = kX$, a proportional relationship in the transformed variable X.

v (m/s)	5	10	15	20	25	30
$X = v^2$	25	100	225	400	625	900
d_b (m)	2.0	7.6	17.9	30.1	49.2	68.5

Using the least-squares formula,

$$k = \frac{\sum X_i d_i}{\sum X_i^2} = \frac{109277.5}{1421875} \approx 0.07685.$$

Therefore

$$\boxed{d_b \approx 0.07685v^2.}$$

Because $d_b = v^2/(2a)$ under constant deceleration, the fitted coefficient corresponds to

$$a \approx \frac{1}{2(0.07685)} \approx 6.51 \text{ m/s}^2.$$

The fitted parameter therefore has a physical interpretation rather than being merely a numerical constant.

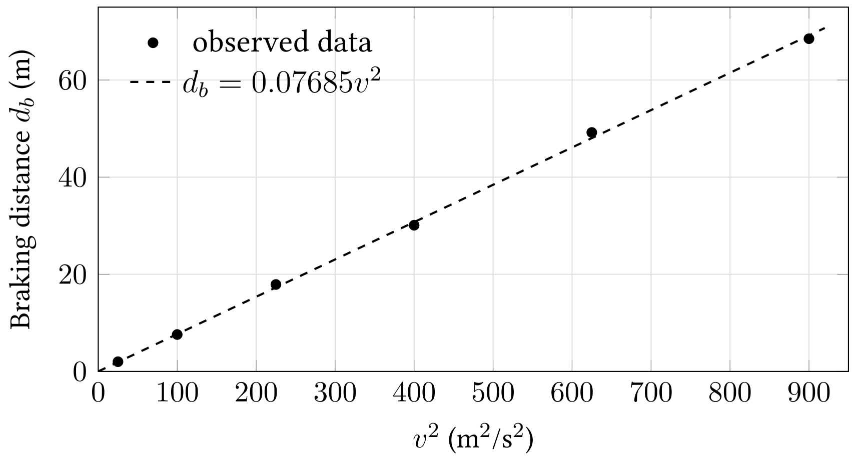

Figure 3: Plotting d_b against v^2 turns the proposed quadratic relationship into a line through the origin.

Model criticism: Should the intercept be forced to zero?

For $v = 0$, physical reasoning requires $d_b = 0$. Forcing the line through the origin is therefore justified. In another problem, a nonzero baseline may be real: an instrument can have an offset, or a delivery fee can include a fixed charge. The fitting constraint must come from the context, not from convenience.

Residual check **[8 min]**

For $v = 25$ m/s, the fitted model predicts

$$\hat{d}_b = 0.07685(25)^2 \approx 48.03 \text{ m}.$$

The observation is 49.2 m.

1. Calculate the residual.
2. Explain whether the model underpredicts or overpredicts.
3. Suggest one physical reason for the discrepancy.

4 Fitting a general straight line

When a nonzero intercept is plausible, fit

$$y = ax + b.$$

The least-squares line minimises

$$S(a,b)=\sum_{i=1}^{m}(y_i-ax_i-b)^2.$$

A convenient form of the solution is

$$\boxed{a=\frac{\sum(x_i-\bar{x})(y_i-\bar{y})}{\sum(x_i-\bar{x})^2},\qquad b=\bar{y}-a\bar{x}.}$$

Worked Example 2: Calibrating a force sensor

Known masses x in kg are placed on a sensor, producing output voltages y.

x (kg)	0	1	2	3	4	5
y (V)	0.18	1.72	3.19	4.76	6.25	7.81

Here $\bar{x}=2.5$, $\bar{y}=3.985$,

$$\sum(x_i-\bar{x})(y_i-\bar{y})=26.655,\qquad \sum(x_i-\bar{x})^2=17.5.$$

Thus

$$a\approx 1.523,\qquad b\approx 0.177,$$

and the calibration model is

$$\boxed{y\approx 1.523x+0.177.}$$

The intercept is meaningful: the unloaded sensor reads about 0.18 V rather than zero. Forcing the model through the origin would hide this offset.

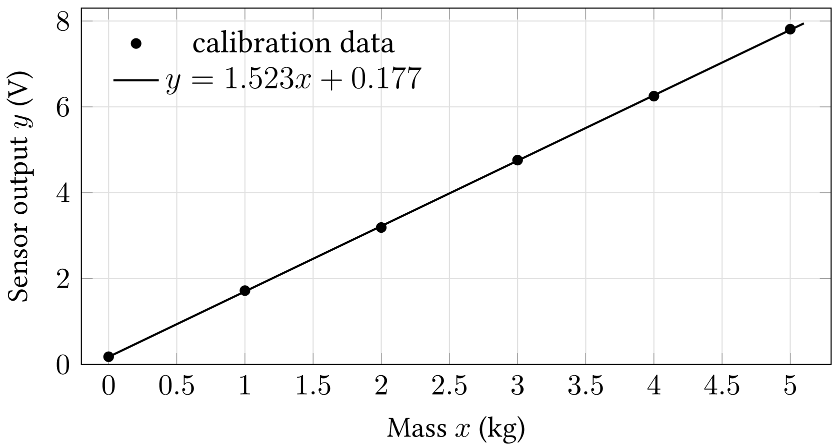

Figure 4: The fitted intercept represents sensor offset. A parameter should be interpreted in the original context.

Why squares?

Least squares minimises

$$\text{SSE}=\sum r_i^2,$$

the **sum of squared errors**. Squaring prevents positive and negative residuals from cancelling and penalises large residuals strongly. This makes the method mathematically convenient, but also sensitive to **outliers**.

Kenneth Cheng · skcKenneth.github.io

5 Error measures and residual plots

No single error measure answers every question. Three common summaries are

$$\text{MAE} = \frac{1}{m}\sum |r_i|, \qquad \text{RMSE} = \sqrt{\frac{1}{m}\sum r_i^2}, \qquad \max |r_i|.$$

The **mean absolute error** is easy to interpret; the **root mean squared error** penalises large errors more strongly; the maximum error focuses on the worst case. All three have the same units as the response variable.

For the braking-distance fit,

$$\text{MAE} \approx 0.54 \text{ m}, \qquad \text{RMSE} \approx 0.66 \text{ m}, \qquad \max |r_i| \approx 1.17 \text{ m}.$$

These values describe the observed calibration range only. They do not prove that the model remains accurate at higher speeds or on different roads.

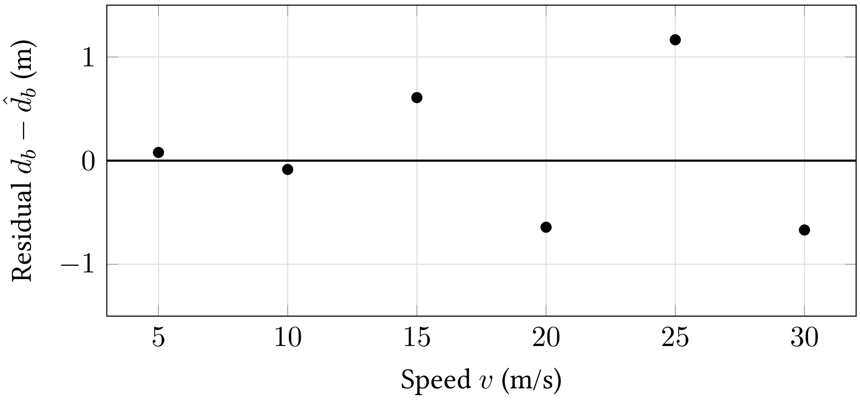

Figure 5: Residuals should be examined for systematic patterns, not merely averaged into one statistic.

Checkpoint: Reading a residual plot

Suppose residuals are positive at low x, negative in the middle, and positive again at high x.

1. Is the scatter random?
2. What does the curved pattern suggest about a fitted straight line?
3. Would collecting more decimal places solve the main problem?

6 Linearising an exponential model

Some model families become linear after a transformation. For an exponential decay model

$$C = Ae^{-kt},$$

taking natural logarithms gives

$$\ln C = \ln A - kt.$$

Thus a plot of $\ln C$ against t should be approximately linear, with slope $-k$ and intercept $\ln A$.

Worked Example 3: Battery capacity decay

A battery is repeatedly tested under the same load. Capacity C is recorded as a percentage of its initial value.

Cycle block t	0	1	2	3	4	5
Capacity C (%)	100	82	68	55	46	37

A least-squares line fitted to $(t, \ln C)$ is approximately

$$\ln C = 4.607 - 0.19765t.$$

Exponentiating gives

$$\boxed{C \approx 100.2e^{-0.19765t}.}$$

The fitted model predicts that capacity is multiplied by

$$e^{-0.19765} \approx 0.821$$

per cycle block, corresponding to an average decrease of about 17.9% of the remaining capacity each block.

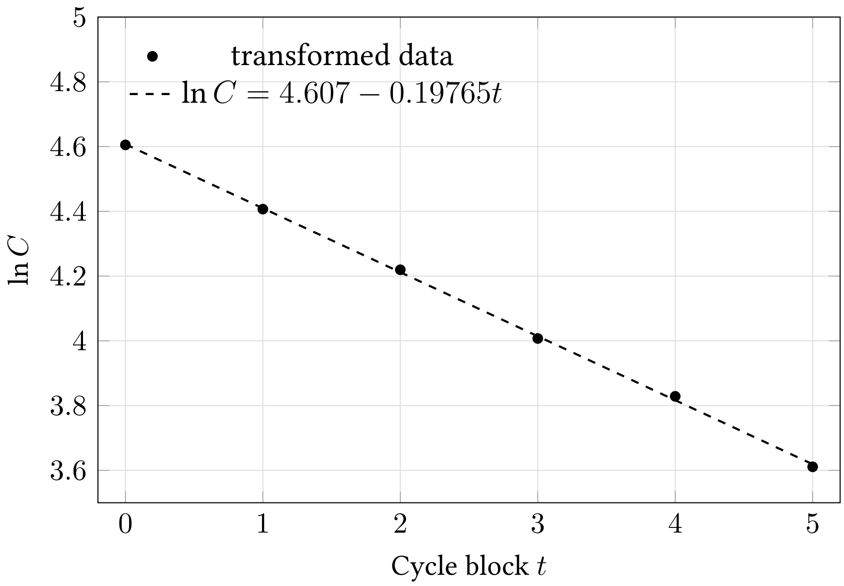

Figure 6: Linearisation makes parameter estimation convenient, but the resulting model must be checked in the original capacity scale.

Model criticism: Transformation changes the fitting criterion

Least squares in $\ln C$ minimises squared differences between logarithms, not squared differences between capacities. It gives greater relative importance to smaller values. Therefore:

1. transform only when the model form and data permit it;
2. convert the fitted model back to the original variables;
3. calculate residuals and error measures in the original scale;
4. compare alternative models using the same response variable.

Choose a transformation **[7 min]**

State the variables that should be plotted to test each proposed model.

1. $y = Ae^{bx}$
2. $y = Ax^n$
3. $y = A/x$

For each case, identify the slope or intercept that has a parameter interpretation.

7 Choosing between models: fit, structure, and validation

A model should not be selected only because it has the smallest calibration error. We also ask whether it respects the mechanism, whether residuals show structure, and whether it predicts data that were not used to fit the parameters.

Worked Example 4: Linear or exponential decay?

Using the six battery observations from $t = 0$ to $t = 5$:

- the exponential fit is $C \approx 100.2e^{-0.19765t}$ with calibration RMSE about 0.39 percentage points;
- the least-squares straight line is $C \approx 95.81 - 12.46t$ with calibration RMSE about 2.93 percentage points.

The value at $t = 6$, $C = 30.5\%$, is withheld during fitting and used for **validation**.

$$\hat{C}_{\text{exp}}(6) \approx 30.61, \qquad \hat{C}_{\text{lin}}(6) \approx 21.07.$$

The exponential model performs much better on the withheld point and never predicts negative capacity. Both the numerical evidence and the mechanism support exponential decay over this range.

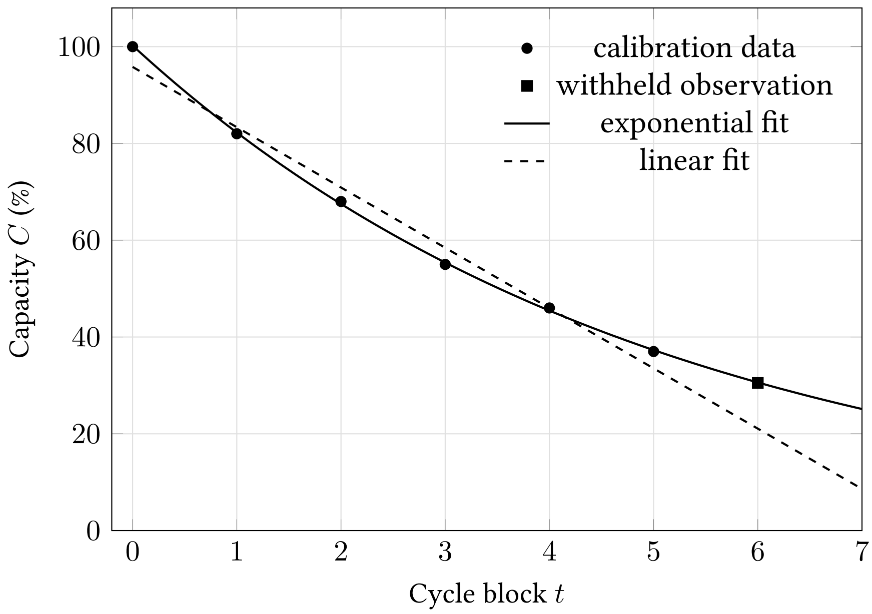

Figure 7: Calibration error, validation performance, and long-term plausibility all favour the exponential model.

A practical model-selection checklist

Before accepting a fitted model, ask:

1. **Mechanism:** Does the functional form agree with the assumptions or governing relationships?
2. **Scale:** Were errors assessed in the original response variable?
3. **Residuals:** Are they small and free of systematic pattern?
4. **Validation**: Does the model predict observations not used for fitting?
5. **Range:** Is the intended prediction inside or outside the observed range?
6. **Interpretability:** Do the parameters and conclusions make sense in context?

Competition-style communication **[10 min]**

Write a four-sentence model-selection paragraph for the battery example:

1. identify the two models compared;
2. report one calibration measure;
3. explain the withheld-data result;
4. state one limitation or condition on the conclusion.

Avoid saying that a high-quality fit "proves" the model.

Optional extension: Modern audit: measurement uncertainty and identifiability

A small residual does not show that every fitted parameter is well determined. If x_i is measured with error but treated as exact, an ordinary least-squares slope can be biased; residuals then combine response noise, input uncertainty, model-form error, and parameter uncertainty. Record instrument resolution and repeated-measurement variability, then propagate plausible perturbations through the fit.

For a model $f(x;\boldsymbol{\theta})$, the parameters are practically **identifiable** only when the available observations can distinguish their effects. A useful local diagnostic is the sensitivity matrix

$$J_{ij} = \frac{\partial f(x_i;\boldsymbol{\theta})}{\partial \theta_j}.$$

If two columns of J are nearly proportional, different parameter combinations produce almost the same fitted curve. More decimal places or a lower training error will not resolve that ambiguity. Report parameter intervals or profiles, test the fit under stated measurement uncertainty, and collect observations where the competing parameter effects differ most. Prediction may still be stable even when individual parameters are not, so audit parameter claims and prediction claims separately.

8 Lesson summary

Seven ideas to retain

1. Model fitting estimates parameters within a proposed model family; interpolation is a different task.
2. A residual is observed minus predicted, with the same units as the response variable.

3. The fitting constraint—through the origin or with an intercept—must be justified by context.
4. Least squares minimises squared residuals and is sensitive to large errors and outliers.
5. Residual plots reveal systematic structure that a single error statistic can hide.
6. Transformations simplify fitting but alter the error criterion; always return to the original variables.
7. A credible model combines mechanism, calibration, validation, interpretation, and stated limitations.

Checkpoint: Exit ticket **[2 min]**

Complete both statements:

1. "A small RMSE is not sufficient because ________________________________."
2. "I would force a fitted line through the origin only when ________________________________."

9 Practice problems

Core practice

1. A model predicts water demand of 82, 91, and 106 L, while the corresponding observations are 85, 88, and 110 L.
   (a) Calculate the three residuals.
   (b) Find the MAE and RMSE.
   (c) Which observation contributes most strongly to the SSE?
2. The data below relate force F to extension x for a spring:

x (cm)	1	2	3	4	5
F (N)	1.9	4.2	6.0	8.1	10.3

   Fit $F = kx$ by least squares. Interpret the units of k and calculate the residual at $x = 4$ cm.
3. A temperature sensor gives readings

x (actual °C)	0	10	20	30	40
y (sensor °C)	1.2	10.8	20.5	30.6	40.1

   Explain why an intercept should be allowed. Fit a straight line using a spreadsheet or calculator and interpret both parameters.
4. State the transformation for each model and explain what a straight transformed plot would support:
   (a) $y = Ae^{bx}$;
   (b) $y = Ax^n$;
   (c) $y = A/(B + x)$.
5. Two models have the following residuals, measured in minutes:

$$A : (-1.0, -0.8, 0.2, 0.4, 1.2), \qquad B : (-0.2, -0.1, 0.0, 0.1, 3.0).$$

Calculate and compare their MAE, RMSE, and maximum error. Which model is preferred by MAE, which by RMSE, and which for avoiding a severe worst case? Explain why "average performance" needs a stated loss measure here.

6. A fitted quadratic has residuals that alternate randomly around zero, while a fitted line has residuals forming a clear U-shape. Explain what this suggests. Why is "the quadratic has more parameters" not by itself a sufficient conclusion?

7. A population model is fitted using years 2018–2024 and then used to predict 2045. Explain why a small training RMSE does not establish that the 2045 forecast is trustworthy. Name two additional checks.

Practice check and hints

1. Residuals are $(3, -3, 4)$ L; MAE $= 3.33$ L, RMSE ≈ 3.37 L, and the third observation contributes most to SSE.
2. $k = 2.04$ N/cm; at $x = 4$ cm the residual is $8.1 - 8.16 = -0.06$ N.
3. A free-intercept fit is $y \approx 0.976x + 1.12$; the intercept represents the zero-load offset.
4. Use $(x, \ln y)$, $(\ln x, \ln y)$, and $(x, 1/y)$, respectively; a line supports the stated family but does not by itself validate it.
5. For A, MAE $= 0.72$, RMSE ≈ 0.81, maximum error $= 1.2$; for B, the values are 0.68, 1.35, and 3.0. Thus MAE prefers B, whereas RMSE and worst-case control prefer A.
6. The U-shaped line residuals indicate missed curvature. Compare withheld performance and residual structure before accepting the quadratic.
7. A credible long-horizon check includes temporal holdout testing plus at least one of mechanism review, scenario sensitivity, or comparison with an alternative model family.

Modeling challenge

Choose one dataset that your team can collect or obtain: cooling temperature, queue waiting time, plant growth, device battery decay, transport travel time, or another approved context.
Prepare a one-page fitting brief containing:

1. the question, variables, units, and proposed model family;
2. a justification for allowing or excluding an intercept;
3. fitted parameter values and one graph in the original variables;
4. a residual plot and at least two error measures;
5. one validation strategy and one limitation;
6. a short conclusion written for a non-specialist decision maker.

The goal is not to obtain the lowest possible error by adding terms. The goal is to produce a defensible model.

A Optional derivation: the least-squares line

This appendix is not part of the two-hour core lesson. It may be assigned to students who are comfortable with partial differentiation.

For $y = ax + b$, minimise

$$S(a,b) = \sum_{i=1}^{m} (y_i - ax_i - b)^2.$$

Setting both partial derivatives equal to zero gives

$$\frac{\partial S}{\partial a} = -2\sum x_i(y_i - ax_i - b) = 0,$$
$$\frac{\partial S}{\partial b} = -2\sum (y_i - ax_i - b) = 0.$$

Therefore

$$a\sum x_i^2 + b\sum x_i = \sum x_i y_i,$$
$$a\sum x_i + mb = \sum y_i.$$

These are the **normal equations**. Solving them is equivalent to the mean-centred formulas used in the lesson.

B Optional extension: alternatives to least squares

Least squares is not the only defensible fitting criterion.

- The L^1 criterion minimises $\sum |r_i|$ and is less sensitive to a single large outlier.
- The L^∞ or **minimax criterion** minimises $\max |r_i|$ and is useful when the worst error matters most.

The appropriate criterion depends on the decision. A medical safety model, a forecasting model, and an engineering calibration problem may not value errors in the same way.

Model audit card

Audit before trusting or communicating a model

Purpose	What question or decision does the model support?	**Data**	Which inputs are observed, derived, or simulated?
Assumptions	Which simplifications may change the answer?	**Units**	Are dimensions and scales consistent?
Verification	Were equations, code, constraints, and limiting cases checked?	**Validation**	Is performance tested against independent evidence?
Uncertainty	How do measurement, parameter, solver, and structural errors affect conclusions?	**Ethics**	Who benefits, who bears risk, and what fairness criterion matters?
Reproducibility	Are data provenance, versions, seeds, and transformations recorded?	**Communication**	Are claims conditional, traceable, and decision-relevant?

Further reading

- Engineering Statistics Handbook (NIST / SEMATECH)
- Introductory Statistics 2e (OpenStax)

Sources, provenance, and licence

These notes follow the iterative modeling and assessment perspective in the SIAM/COMAP GAIMME report. Competition-specific remarks are a snapshot checked on 20 August 2026 against the official COMAP HiMCM instructions; readers should verify current rules before submission. Unless a source is explicitly identified, classroom datasets and numerical scenarios are synthetic or simulated teaching examples and are not empirical claims about a real institution. Third-party references retain their own terms. Original teaching material by Kenneth Sok Kin Cheng is released under CC BY-NC-SA 4.0.

Student response guide and checking criteria

For closed calculations, show the definition, substitution, units, and at least one independent check; the worked examples in this unit are method hints, but you must complete the arithmetic yourself. Open modeling responses are checked on five dimensions: a clear purpose and output; stated data and provenance; defensible assumptions and scope; traceable computational verification and model validation; and alignment among uncertainty, limitations, and the recommendation. Support each dimension with concrete evidence rather than writing only "completed."